arXiv:1511.08535v3 [math.GR] 23 Aug 2016

Submitted exclusively to the London Mathematical Society
doi:10.1112/0000/000000

A Diameter Bound for Finite Simple Groups of Large Rank

Arindam Biswas [†], Yilong Yang [‡]

Abstract

Given a non-abelian finite simple group G of Lie type, and an arbitrary symmetric generating set S, it is conjectured by László Babai that its Cayley graph $\Gamma(G, S)$ will have a diameter bound of $(\log |G|)^{O(1)}$. However, little progress has been made when the rank of G is large. In this article, we shall show that if G has rank n, and its base field has bounded size, then the diameter of $\Gamma(G, S)$ would be bounded by $\exp(O(n(\log n)^3))$.

1. *Introduction*

1.1. *History and Background*

Given a group G and a symmetric generating set S, one can construct a corresponding Cayley graph $\Gamma(G, S)$. Its vertices are elements of G, and two vertices $g, h \in G$ are connected by an edge iff there is an element $s \in S$ such that $sg = h$. The Cayley graph is a metric space where the distance between two vertices is simply the length of the shortest path from one to the other. In this way, we can discuss the diameter of the Cayley graph. Equivalently, we can also define the diameter to be the smallest number ℓ such that every element of G can be written as a product of at most ℓ elements of S.

If G is a non-abelian finite simple group, we expect all its Cayley graphs to have good connectivity. In particular, we have the following conjecture of Babai:

Conjectures 1.1 Babai, [BS92]. For any non-abelian finite simple group G, and for any symmetric generating set, the diameter of the Cayley graph is bounded by $(\log |G|)^{O(1)}$, where the implied constant is absolute.

The first class of simple groups verified for Babai's conjecture was $\mathrm{PSL}_2(\mathbb{Z}/p\mathbb{Z})$ with p prime, by Helfgott [Hel08]. Afterwards, a lot of research was done on the diameters and related expansion properties of these Cayley graphs.

The best result to date are those by Pyber and Szabo [PS16], and Breuillard, Green and Tao [BGT11a]. They verified Babai's conjecture for all finite simple groups of Lie type with bounded rank.[§]

For all non-abelian finite simple groups, Breuillard and Tointon [BT16] also obtained a diameter bound of $\max(|G|^{\epsilon}, C_{\epsilon})$ for arbitrary $\epsilon > 0$ and a constant C_{ϵ} depending only on ϵ. The diameter bounds in all these previous results depend poorly on the rank of the group. It

[†]Laboratoire de Mathématiques, Bâtiment 425, Faculté des Sciences d'Orsay, Université Paris-Sud XI, 91405 Orsay Cedex, France. Email: arindam.biswas@math.u-psud.fr.

[‡]3170 Sawtelle Blvd Apt 203, Los Angeles, CA 90066, USA. Email: yy26@math.ucla.edu

2000 *Mathematics Subject Classification* 20G40, 20F69 (primary), 05C25, 20E32, 51N30, 11N13, 05C81 (secondary).

[§]The preprint of PS was published on arXiv in 2010, and proved Babai's conjecture for all finite simple groups of Lie type with bounded ranks. Unlike PS, BGT first announced these results only for finite simple groups not belonging to the Suzuki and Ree family. Their method also applies to these cases, but this only appeared later in [BGT11b] and [Bre+15]. We thank László Pyber for this remark.

is one of the aims of this paper to improve the dependency on the rank, in the case of finite simple groups of Lie type.

On the other hand, a lot of research was also done for the symmetric group S_n and the alternating group A_n. In 1988, Babai and Seress showed the following theorem.

THEOREM 1.2 Babai and Seress, [BS88]. *Let $G = \mathrm{S}_n$ or $G = \mathrm{A}_n$, then for any symmetric generating set,*

$$\mathrm{diam}(G) \leq \exp(\sqrt{n \log n}(1 + o_n(1)) = \exp(\sqrt{\log |G|}(1 + o_n(1)).$$

This was the best known bound for S_n or A_n for over two decades, until Helfgott and Seress recently showed the following.

THEOREM 1.3 Helfgott and Seress, [HS14]. *Let $G = \mathrm{S}_n$ or $G = \mathrm{A}_n$, then for any symmetric generating set,*

$$\mathrm{diam}(G) \leq \exp(O((\log n)^4 \log \log n)) = \exp((\log \log |G|)^{O(1)})$$

.

In this article we give a modest upper bound on the diameter for finite simple groups of Lie type, where the dependency on rank is lessened.

THEOREM 1.4 Main Theorem. *Let G be a finite simple group of Lie type, with rank n and base field $\mathbb{F}_q$, then*

$$\mathrm{diam}(G) \leq q^{O(n((\log n + \log q)^3))}.$$

In particular, if the base field has bounded size, we have

$$\mathrm{diam}(G) \leq \exp(O(n(\log n)^3)) = \exp(O(\sqrt{\log |G|}(\log \log |G|)^3)).$$

1.2. *Preliminaries*

DEFINITION 1.5. Given an algebraic group G over a finite field $\mathbb{F}_q$, the ***algebraic rank*** of G is the dimension of a maximal torus in G.

PROPOSITION 1.6. *There is an absolute constant C, such that any finite simple group of Lie type G of algebraic rank larger than C must be a projective special linear group $\mathrm{PSL}_n(\mathbb{F}_q)$, a projective symplectic group $\mathrm{PSp}_n(\mathbb{F}_q)$, a projective special unitary group $\mathrm{PSU}_n(\mathbb{F}_q)$, or the simple quotient $\mathrm{P}\Omega_n(\mathbb{F}_q)$ of the derived subgroup $\Omega_n(\mathbb{F}_q)$ of the orthogonal group $\mathrm{O}_n(\mathbb{F}_q)$.*

Proof. Going through the list of finite simple groups of Lie type, the algebraic ranks of all but the above four families of groups are bounded by an absolute constant C. □

In this paper, we are only interested in finite simple groups of Lie type with large ranks. Therefore, these four families of groups are all we need to deal with.

PROPOSITION 1.7. *The algebraic rank of groups $\mathrm{GL}_n(\mathbb{F}_q)$, $\mathrm{SL}_n(\mathbb{F}_q)$, $\mathrm{PSL}_n(\mathbb{F}_q)$, $\mathrm{Sp}_n(\mathbb{F}_q)$, $\mathrm{PSp}_n(\mathbb{F}_q)$, $\mathrm{U}_n(\mathbb{F}_q)$, $\mathrm{SU}_n(\mathbb{F}_q)$, $\mathrm{PSU}_n(\mathbb{F}_q)$, $\mathrm{O}_n(\mathbb{F}_q)$, $\Omega_n(\mathbb{F}_q)$, $\mathrm{P}\Omega_n(\mathbb{F}_q)$ are between $\frac{n}{2} - 1$ and n.*

Proof. This is done by computing their algebraic rank one by one. □

Let G be any group. Given a subset S of G, we shall use S^d to denote $\{g_1...g_d : g_1, ..., g_d \in S\}$, i.e., the set of product of d elements of S.

We shall use the "big O" and "small o" convention. We let $O(g(x))$ to denote a quantity whose absolute value is bounded by $Cg(x)$ for some absolute constant C. And we let $o_n(1)$ to denote a quantity whose limit is 0 as n goes to infinity.

1.3. *Outline of the Proof*

DEFINITION 1.8. Given a matrix A, we define its ***degree*** to be $\deg(A) = \mathrm{rank}(A - I)$. Equivalently, it is n minus the dimension of the fixed subspace of A.

The general idea is to find a matrix of small degree and close to the identity. And from there on, the expansion would be very fast. This is analogous to [BS88], the case of symmetric and alternating groups, where one first finds elements of small support (i.e., 2-cycles or 3-cycles), and then proceed to fill up the whole group.

Given any finite simple group of Lie type, we proceed via the following steps. Let G be $\mathrm{SL}_n(\mathbb{F}_q)$, $\mathrm{Sp}_n(\mathbb{F}_q)$, $\mathrm{SU}_n(\mathbb{F}_q)$ or $\Omega_n(\mathbb{F}_q)$. Then we can assume that G acts on a vector space V of dimension n.

(1) We start by finding a subset of small diameter in G near the identity, which we call a t-transversal set. Such a set contains an extension for every linear or isometric embeddings of any t-dimensional subspace W into V. This is an analogy to a t-transitive subset of a symmetric group S_n. See Section 2 for the linear case, and Section 6 for the symplectic case, the unitary case, and the orthogonal case.

(2) Using the t-transversal set above, we can find a special matrix, which we call a P-matrix. Some large power of a P-matrix will have a much smaller degree than the original one. This is the main degree reducing step, inspired by Lemma 1 of the paper of Babai and Seress [BS88]. The result on P-matrices is dependent on an inequality of primes, which is deduced in Section 3. For P-matrices, see Section 4.

(3) Combined with repeated commutators with carefully chosen elements from the t-transversal set, we can repeat the above step many times, until we reach a matrix of very small degree. See Section 5 for the linear case, and Section 7 for the symplectic case, the unitary case, and the orthogonal case.

(4) From a small degree matrix, one can quickly fill up its conjugacy class. Then by the result of Liebeck and Shalev [LS01], from a conjugacy class, one can quickly fill up the whole group. See Section 8.

1.4. *Acknowledgements*

We would like to thank Emmanuel Breuillard and Terence Tao for a number of helpful discussions and advice in this subject.

2. *t-Transversal Sets with Small Diameters*

Let V be a vector space of dimension n over the field $\mathbb{F}_q$. Let the group $\mathrm{GL}_n(\mathbb{F}_q)$ act on it naturally.

DEFINITION 2.1. A subset S of $\mathrm{GL}_n(\mathbb{F}_q)$ is called a ***t-transversal set*** if given any embedding X of a t-dimensional subspace W into V, we can find $A \in S$ that extends X on W.

LEMMA 2.2. *$\mathrm{GL}_n(\mathbb{F}_q)$ is t-transversal for all t, and $\mathrm{SL}_n(\mathbb{F}_q)$ is t-transversal for all $t < n$.*

Proof. Let W be any subspace with a basis $w_1, ..., w_t$. We can complete this into a basis of V with new vectors $v_1, ..., v_{n-t}$. Let A be a matrix with column vectors $w_1, ..., w_t, v_1, ..., v_{n-t}$. In the case when $t < n$, we can multiply v_{n-t} by a constant so that $\det(A) = 1$.

For any embedding X of W into V, $X(w_1), ..., X(w_t)$ are linearly independent. We can complete this into a basis of V with new vectors $u_1, ..., u_{n-t}$. Let B be a matrix with column vectors $X(w_1), ..., X(w_t), u_1, ..., u_{n-t}$. In the case when $t < n$, we can multiply u_{n-t} by a constant so that $\det(B) = 1$.

Now $B(A)^{-1}$ is in $\mathrm{GL}_n(\mathbb{F}_q)$ and, if $t < n$, also in $\mathrm{SL}_n(\mathbb{F}_q)$. We also have $(B(A)^{-1})|_W = X$. □

LEMMA 2.3. *For any symmetric subset S of $\mathrm{GL}_n(\mathbb{F}_q)$, if the subgroup generated by S is t-transversal, then $\bigcup_{d=1}^{d=q^{nt}} S^d$ is t-transversal.*

Proof. Let W be any t-dimensional subspace. Let $L(W)$ be the set of embeddings of W into V. Let H be the subgroup generated by S. Then an element g of H acts on $L(W)$ by $g(X) = (g \circ X)|_W$ for any $X \in L(W)$. Let Γ be the corresponding Schreier graph of this action of H on $L(W)$ with generating set S, i.e., the vertices are elements of $L(W)$, and two vertices X, Y are connected iff $g(X) = Y$ for some $g \in S$.

Now, since H is t-transversal, the graph Γ is connected. So the diameter of Γ is trivially bounded by its number of vertices, which is at most q^{nt}. As a result, the set $\bigcup_{d=1}^{q^{nt}} S^d$ is t-transversal. □

COROLLARY 2.4. *Given any symmetric generating set S for $\mathrm{GL}_n(\mathbb{F}_q)$, the set $\bigcup_{d=1}^{d=q^{nt}} S^d$ is t-transversal. If $t < n$, then the same statement is true with $\mathrm{SL}_n(\mathbb{F}_q)$ replacing $\mathrm{GL}_n(\mathbb{F}_q)$.*

3. *An Inequality on Primes*

In this section, we shall establish an inequality on primes to be used in the next section.

Throughout this section, we shall fix a prime p_0 and fix a power of it q_0, which in the next section shall become the characteristic and the order of a finite field.

Let $p_1, ..., p_r$ be the first r primes coprime to $p_0(q_0 - 1)$. Let M be the least common multiple of $p_1 - 1, p_2 - 1, p_3 - 1, ..., p_r - 1$. Let S be the sum of $p_1 - 1, p_2 - 1, p_3 - 1, ..., p_r - 1$. Our goal for this section is the following proposition:

LEMMA 3.1. *There exist absolute constants c_1 and c_2 such that, if $p_r \geq c_1 \log q_0$, then*

$$S \leq (p_r)^2 \leq c_2(\log M)^3.$$

Before we prove this, let us first set up more notations. Let P^+ be the function that sends each positive integer to its largest prime factor. Let $P = \{p_1, ..., p_r\}$. For any $\delta > 0$, let $P_\delta = \{\text{prime number } p : 3 \leq p \leq p_r, P^+(p-1) \geq (p_r)^\delta\}$, and let $P_\delta^* = P_\delta \cap P$.

We start by citing an important theorem of Fouvry.

LEMMA 3.2 (Fouvry [Fou85]). *There is an absolute constant $\delta > \frac{2}{3}$, and an absolute constant c_0, c_p, such that for $p_r \geq c_p$, we have*

$$|P_\delta| \geq c_0 \frac{p_r}{\log p_r}.$$

COROLLARY 3.3. *There is an absolute constant $\delta > \frac{2}{3}$, and absolute constants c_0 and c_3, such that $|P_\delta^*| \geq c_0 \frac{p_r}{\log p_r} - c_3 \frac{\log q_0}{\log\log q_0}$ if $q_0 > e^e$, and $|P_\delta^*| \geq c_0 \frac{p_r}{\log p_r} - 3$ if $q_0 \leq e^e$.*

Proof. By prime number theorem, when $\log\log q_0 > 1$ (i.e., $q_0 > e^e \approx 15.15$), the number of prime factors of $p_0(q_0 - 1)$ is bounded by $c_3 \frac{\log q_0}{\log\log q_0}$ for some absolute constant c_3. If $q_0 \leq e^e$, then there are at most 3 prime factors of $p_0(q_0 - 1)$. □

LEMMA 3.4. *Let $p \geq (p_r)^\delta$ be some prime. Then*

$$|((P^+)^{-1}(p) + 1) \cap P_\delta^*| \leq \frac{2p_r}{(\log 2)(p_r^\delta - 1)}.$$

Proof.

We assume that the left hand side of the inequality is non-zero, because otherwise the inequality is trivial. Then p divides some $p_i - 1$, which is not a prime. So in particular, $2p < p_r$.

The set $((P^+)^{-1}(p) + 1) \cap P_\delta^*$ is contained in the set of primes $\leq p_r$ that are congruent to 1 mod p. By the Brun-Titchmarsh theorem, combined with the fact that $p \geq (p_r)^\delta$, we have

$$\begin{aligned}
|(P^+)^{-1}(p) \cap P_\delta^*| \leq & \frac{2p_r}{\phi(p)\log\frac{p_r}{p}} \\
\leq & \frac{2p_r}{(\log 2)(p-1)} \\
\leq & \frac{2p_r}{(\log 2)((p_r)^\delta - 1)}.
\end{aligned}$$

Here ϕ is the Euler totient function. □

Now we have enough to prove Lemma 3.1.

Proof of Lemma 3.1. The first inequality is straightforward

$$S \leq \sum_{\text{prime } p \leq p_r} p \leq (p_r)^2.$$

All the primes in $P^+(P_\delta^* - 1)$ are factors of M, and they are all larger than $(p_r)^\delta$. Furthermore, when $q_0 > e^e$, we have

$$\begin{aligned}
|P^+(P_\delta^* - 1)| \geq & \frac{|P_\delta^*|}{\max_{p \geq (p_r)^\delta} |((P^+)^{-1}(p) + 1) \cap P_\delta^*|} \\
\geq & (c_0 \frac{p_r}{\log p_r} - c_3 \frac{\log q_0}{\log\log q_0}) / (\frac{2p_r}{(\log 2)((p_r)^\delta - 1)}) \\
\geq & \frac{\log 2}{2}((p_r)^\delta - 1)(\frac{c_0}{\log p_r} - \frac{c_3}{p_r}\frac{\log q_0}{\log\log q_0}).
\end{aligned}$$

So, if $p_r > c_1 \log q_0$ for an absolute constant c_1 such that $c_3 \frac{1+\log c_1}{c_1} < \frac{c_0}{2}$, then we have

$$\begin{aligned}\log M \geq & |P^+(P_\delta^* - 1)| \log((p_r)^\delta) \\ \geq & \frac{\log 2}{2}\delta((p_r)^\delta - 1)(c_0 - c_3 \frac{\log q_0}{\log\log q_0}\frac{\log p_r}{p_r}) \\ \geq & \frac{\log 2}{2}\delta((p_r)^\delta - 1)(c_0 - c_3 \frac{1+\log c_1}{c_1}) \\ \geq & \frac{\log 2}{4}\delta c_0((p_r)^\delta - 1).\end{aligned}$$

Since $\delta > \frac{2}{3}$, we can pick some constant such that $c_2(\log M)^3 \geq (p_r)^2$.
Now, suppose $q_0 \leq e^e$. Then similarly we have

$$\begin{aligned}|P^+(P_\delta^* - 1)| \geq & \frac{|P_\delta^*|}{\max_{p \geq (p_r)^\delta} |((P^+)^{-1}(p) + 1) \cap P_\delta^*|} \\ \geq & (c_0 \frac{p_r}{\log p_r} - 3)/(\frac{2p_r}{(\log 2)((p_r)^\delta - 1)}) \\ \geq & \frac{\log 2}{2}((p_r)^\delta - 1)(\frac{c_0}{\log p_r} - \frac{3}{p_r}).\end{aligned}$$

So, if $p_r > c_1 \log q_0$ for a sufficiently large absolute constant c_1 such that $\frac{3\log p_r}{p_r} < \frac{c_0}{2}$, then we have

$$\begin{aligned}\log M \geq & |P^+(P_\delta^* - 1)| \log((p_r)^\delta) \\ \geq & \frac{\log 2}{2}\delta((p_r)^\delta - 1)(c_0 - \frac{3\log p_r}{p_r}) \\ \geq & \frac{\log 2}{4}\delta c_0((p_r)^\delta - 1).\end{aligned}$$

Since $\delta > \frac{2}{3}$, we can again pick some constant such that $c_2(\log M)^3 \geq (p_r)^2$.

□

As a side note, for any improved value of δ in the Fouvry's theorem, our diameter bound in this paper would improve to $q^{O(n(\log n + \log q)^{\frac{2}{\delta}})}$ for finite simple groups of Lie type of rank n over $\mathbb{F}_q$.

If one were to assume the Hardy-Littlewood conjucture on prime tuples, the δ could be improved to $1 - o(1)$. Combine this with the more efficient estimate $S \leq \frac{(p_r)^2}{\log p_r}$, the diameter bound of this paper would improve to $q^{O(n(\log n + \log q)^2)}$ for finite simple groups of Lie type of rank n over $\mathbb{F}_q$.

4. *P-Matrices and Degree Reduction*

This section aims to show that, given a P-matrix, we can reduce its degree by raising it to a large power.

DEFINITION 4.1. Let $\mathbb{F}_q$ be a finite field of characteristic p, and let $p_1, p_2, ..., p_r$ be the first r primes coprime to $p(q-1)$. Then a matrix A over $\mathbb{F}_q$ is called a ***P(r)-matrix*** if, for each $i \leq r$, it has a primitive p_i-th root of unity in the algebraic closure of $\mathbb{F}_q$ as an eigenvalue.

LEMMA 4.2. *Let A be a matrix over $\mathbb{F}_q$, a field with characteristic p. Let m be any number coprime to p. Then if A has a primitive m-th root of unity as an eigenvalue, A must have degree at least $o_m(q)$. Here $o_m(q)$ denotes the multiplicative order of p in $\mathbb{Z}/m\mathbb{Z}$.*

Proof. Let $\Phi_m(X)$ be the m-th cyclotomic polynomial. Then by Galois theory over finite field, the polynomial $\Phi_m(X)$ factors into distinct irreducible polynomials of degree $o_m(q)$. See, e.g., [Lan02].

Therefore, if A has a primitive m-th root of unity as an eigenvalue, then A must have at least $o_m(q)$ primitive m-th roots of unity as eigenvalues, and the result follows. □

LEMMA 4.3. *Let A be a non-identity matrix over finite field $\mathbb{F}_q$ of characteristic p. Suppose A has a primitive m-th root of unity as an eigenvalue in the algebraic closure of $\mathbb{F}_q$. Let $P(m)$ be the set of all prime divisors of m coprime to $p(q-1)$. Then A has degree at least $\mathrm{lcm}_{x\in P(m)}(x-1)$. Here* lcm *denote the least common multiple.*

Proof. By replacing m by its factor, we can assume that m is the product of primes in $P(m)$. In particular, we can assume that m is coprime with p. Then A has degree at least $o_m(q)$.

For each $x \in P(m)$, since x is a prime, $o_x(q) = x-1$. Since $o_x(q)$ divides $o_m(q)$ whenever x divides m, therefore $x-1$ is a factor of $o_m(q)$ for all $x \in P(m)$.

So the degree of A is at least $o_m(q) \geq \mathrm{lcm}_{x\in P(m)}(x-1)$. □

LEMMA 4.4. *Let n be an integer, and let q be a power of the prime p. Then we can find an integer r and an absolute constant c, such that the following is true:*

(i) *If $p_1, p_2, ..., p_r$ are the first r primes coprime to $p(q-1)$, then $\mathrm{lcm}_{i=1}^{r}(p_i-1) > n^4$, and $\sum_{i=1}^{r}(p_i-1) < c(\log n + \log q)^3$.*
(ii) *Let $A \in \mathrm{GL}_n(\mathbb{F}_q)$ where the field has characteristic p, and $\deg A = k$. If A is a $P(r)$-matrix, then there exists $\ell \in \mathbb{N}$ such that A^ℓ will be a non-identity matrix of degree at most $\frac{k}{4}$, and 1 is the only eigenvalue of A^ℓ lying in $\mathbb{F}_q$.*

Proof.

The First Statement:

Let M be the least common multiple, and let S be the sum. Let c_1 be the constant as in Lemma 3.1.

Pick p_r to be the smallest prime such that $M > n^4$ and $p_r > c_1 \log q$. Then the second condition guarantees that $S < c_2(\log M)^3$, according to Lemma 3.1.

Now, if $p_r \leq 2c_1 \log q$, then for some absolute constant c_4 by the Prime Number Theorem, we have

$$\begin{aligned}\log M \leq & \sum_{i=1}^{r} \log p_r \\ \leq & c_4 p_r \\ \leq & 2c_1 c_4 \log q.\end{aligned}$$

So $S \leq c(\log q)^3$ for some absolute constant c.

Suppose $p_r > 2c_1 \log q$. Then by the Bertrand-Chebyshev Theorem, $p_{r-1} > c_1 \log q$. Let $M' = \mathrm{lcm}_{i=1}^{r-1}(p_i - 1)$. Then by the minimality of p_r, we must have $M' \leq n^4$. In particular, we have

$$\begin{aligned} 4\log n \geq & \log M' \\ \geq & (\frac{p_{r-1}}{\sqrt{c_2}})^{\frac{2}{3}}. \end{aligned}$$

So, we have $p_{r-1} \leq 8\sqrt{c_2}(\log n)^{\frac{3}{2}}$. Then $p_r \leq 16\sqrt{c_2}(\log n)^{\frac{3}{2}}$.

Furthermore, we have

$$\begin{aligned} \log M \leq & \log(M' p_r) \\ \leq & \log(n^4(16\sqrt{c_2}(\log n)^{\frac{3}{2}})) \\ < & 6\log n + \log(16\sqrt{c_2}). \end{aligned}$$

So $S < c_1(\log M)^3 < c(\log n)^3$ for some absolute constant c.

The Second Statement:

Let M_i denote the least common multiple of $p_1 - 1, p_2 - 1, ..., p_i - 1$. Let $t_1 = p_1 - 1$ and $t_i = \frac{M_i}{M_{i-1}}$ for $i > 1$. Then $\prod_{i=1}^{r} t_i = M_r > n^4$.

Let $N = \{1, 2, ..., n\}$. Let $d_1, ..., d_n$ be the eigenvalues of A in the algebraic closure of $\mathbb{F}_q$.

For each $j \in N$, let P_j be the set of prime factors of the multiplicative order of d_j among $p_1, ..., p_r$. Then by Lemma 4.3, for each $j \in N$,

$$\prod_{p_i \in P_j} t_i \leq \mathrm{lcm}_{p_i \in P}(p_i - 1) \leq k.$$

Now let $n(i)$ denote the number of P_j that contain p_i.

We take the weighted average T of these $n(i)$ with weight $\log t_i$. The sum of the weights is $\sum_{i=1}^{r} \log t_i > 4\log n$.

$$\begin{aligned} T = & \frac{\sum_{1 \leq i \leq r} n(i) \log t_i}{\sum \log t_i} \\ = & \frac{\sum_{1 \leq i \leq r} \sum_{j \in N} (\log t_i) 1_{p_i \in P_j}}{\sum \log t_i} \\ = & \frac{\sum_{j \in N} \sum_{1 \leq i \leq r} (\log t_i) 1_{p_i \in P_j}}{\sum \log t_i} \\ \leq & \frac{\sum_{j \in N} \sum_{p_i \in P_j} \log t_i}{4\log n} \\ \leq & \frac{k \log k}{4 \log n} \\ \leq & \frac{k}{4}. \end{aligned}$$

So there is a p_i such that $n(i) \leq \frac{k}{4}$. So if A has order $m(A)$, then $A^{\frac{m(A)}{p_i}}$ is the desired non-identity matrix of degree at most $\frac{k}{4}$. Every eigenvalue of the latter matrix not equal to 1 is a primitive p_i-th root of unity, which would be outside of $\mathbb{F}_q$. □

5. *Commutators and Degree Reduction*

In this section, we shall use repeated commutators with elements of a t-transversal set. This way, we repeatedly create P-matrices and raise them to a large power, and would eventually end up with a matrix of very small degree.

DEFINITION 5.1. Given any element g of a group G and a symmetric generating set S for G, the ***length*** of g is $\ell(g) = \min\{d \in \mathbb{N} : g = s_1...s_d \text{ for some } s_1, ..., s_d \in S\}$.

PROPOSITION 5.2. *For any matrices A, B, $\deg(ABA^{-1}B^{-1}) \leq 2\min(\deg A, \deg B)$.*

Proof.

$$\begin{aligned}\deg(ABA^{-1}B^{-1}) &= \operatorname{rank}(ABA^{-1}B^{-1} - I)\\ &= \operatorname{rank}(AB - BA)\\ &= \operatorname{rank}((A-I)(B-I) - (B-I)(A-I))\\ &\leq \operatorname{rank}(A-I)(B-I) + \operatorname{rank}(B-I)(A-I)\\ &\leq \operatorname{rank}(A-I) + \operatorname{rank}(A-I)\\ &= 2\operatorname{rank}(A-I).\end{aligned}$$

Similarly, we also have $\deg(ABA^{-1}B^{-1}) \leq 2\operatorname{rank}(B-I)$. So we are done. □

LEMMA 5.3. *Fix any matrix $A \in \mathrm{GL}(V)$ of degree k, such that the eigenvalues of A are either 1 or outside of $\mathbb{F}_q$. For any $t \leq \frac{k}{2}$, we can find a subspace W of V with the following properties:*

(i) $\dim W = t$;
(ii) $W \cap AW = \{0\}$.

Proof. We shall proceed by induction on the dimension of W. Let V_A be the subspace of fixed points of A in V.

Initial Step: Suppose $t = 1$. Simply pick any vector v outside of V_A, and let W be the span of v. We have $W \cap V_A = \{0\}$ by choice of v. Since A has no eigenvalue in $\mathbb{F}_q$ other than 1, v and Av must be linearly independent. So $W \cap AW = \{0\}$.

Inductive Step: Suppose we have found a subspace W of dimension $t-1$ such that $W \cap AW = \{0\}$. I claim that, when $t \leq \frac{k}{2}$, we can find another vector v, such that the desired subspace is the span of v and W.

To prove the existence of v, let us count the number of vectors to avoid. We want v to avoid $V_A + W + AW$. Afterwards, it is enough to let Av avoid any linear combination of v and $W + AW$. So we need v to avoid $\bigcup_{x\in\mathbb{F}_q}(A-x)^{-1}(W+AW)$. Here we shall interpret $(A-x)^{-1}$ as the pullback map of subsets.

Now, since A has no eigenvalue in $\mathbb{F}_q$ other than 1, therefore $A - x$ is invertible when $x \neq 1$. And $A - 1$ has kernel exactly V_A, which has dimension $n - k$. So, we have

$$\begin{aligned}&|(V_A + W + AW) \cup (\bigcup_{x\in\mathbb{F}_q}(A-x)^{-1}(W+AW))|\\ \leq& q^{n-k+2t-2} + q^{n-k+2t-2} + (q-1)q^{2t-2}\\ <& q^{n-k+2t}.\end{aligned}$$

So as long as $2t \leq k$, we have $q^{n-k+2t} \leq q^n$. So it is possible to choose a vector v as desired. □

LEMMA 5.4. *Let A be a matrix in the group $\mathrm{GL}_n(\mathbb{F}_q)$. Then the order of A is less than q^n.*

Proof. By the rational canonical form of a matrix A (see, e.g., [Her75]), it is enough to prove the case when A is a single rational jordan block.

In this case, the characteristic polynomial $f(x)$ of A is a power of an irreducible polynomial $g(x)$. Say $f = g^t$ and the characteristic of $\mathbb{F}_q$ is p. Then the order of A is $(q^{\frac{n}{t}} - 1)p^{\lceil \frac{\log t}{\log p} \rceil} < q^{\frac{n}{t}} p^{\lceil \frac{\log t}{\log p} \rceil}$.

Then it is enough to show that $p^{\lceil \frac{\log t}{\log p} \rceil} \leq q^{(1-\frac{1}{t})n}$. And since $p^{\lceil \frac{\log t}{\log p} \rceil}$ is the smallest p-power that is larger than or equal to t, and since q is a power of p, it is enough to show that $t \leq q^{(1-\frac{1}{t})n}$. And this is true since $q \geq 2$ and $n \geq t$. □

PROPOSITION 5.5. *For any symmetric generating set S of $\mathrm{GL}_n(\mathbb{F}_q)$, $\mathrm{GL}_n(\mathbb{F}_q)$ has a non-trivial element of degree at most $C(\log n + \log q)^3$ for some absolute constant C, of length less than $q^{C' n (\log n + \log q)^3}$ for some absolute constant C'. The same statement is true with $\mathrm{SL}_n(\mathbb{F}_q)$ replacing $\mathrm{GL}_n(\mathbb{F}_q)$.*

Proof. We pick r and c according to Lemma 4.4. We may assume that $c(\log n + \log q)^3 < n$, because otherwise the statement is trivial.

Let $p_1, ..., p_r$ be the first r primes coprime to $p(q-1)$. Let $f_i(x)$ be the irreducible polynomial over $\mathbb{F}_q$ for all the primitive p_i-th roots of unity, and let C_i be the companion matrix of $f_i(x)$.

Initial Step:

Let us find our first $\mathrm{P}(r)$-matrix. Let T be a $c(\log n + \log q)^3$-transversal set. Then by definition, we can find $A_0 \in T$ that maps some subspace W of dimension $c(\log n + \log q)^3$ onto itself, and that its restriction to this subspace is the matrix $(\bigoplus_{i=1}^{r} C_i) \oplus I$ for some arbitrary choices of basis on W, where I is some identity matrix of suitable size.

In particular, A_0 is a $\mathrm{P}(r)$-matrix. Since $A_0 \in T$, by choosing T as in Corollary 2.4, A_0 has length bounded by $q^{cn(\log n + \log q)^3}$.

By using Lemma 4.4, we can raise A_0 to a large power, and obtain a non-identity matrix A_1 of degree $\leq \frac{\deg(A_0)}{4} \leq \frac{n}{4}$, with eigenvalues either 1 or outside of $\mathbb{F}_q$. Since the order of A_0 is bounded by q^n by Lemma 5.4, the length of A_1 is bounded by $q^{cn(\log n + \log q)^3 + n}$.

Inductive Step:

Suppose we have obtained a non-identity matrix A_j with eigenvalues either 1 or outside of $\mathbb{F}_q$, degree at most $\frac{n}{2^{j+1}}$, and length at most $q^{2cn(\log n + \log q)^3 + j(n+2)}$. If $\deg A_j \leq 2c(\log n + \log q)^3$, then we stop. If not, then let us construct a non-identity matrix A_{j+1} of even smaller degree.

First we shall transform A_j into a $\mathrm{P}(r)$-matrix. Find a subspace W_j of dimension at least $c(\log n + \log q)^3$ as in Lemma 5.3 using A_j. In particular, W_j has trivial intersection with $A_j W_j$. Let T' be a $2c(\log n + \log q)^3$-transversal set, then we can find $M_j \in T'$ that fixes $A_j W_j$, and restricts to a map from W_j to W_j as $\bigoplus_{i=1}^{r} C_i \oplus I$ for an arbitrary basis of W_j and some identity matrix I of suitable size.

Consider the commutator $M_j A_j^{-1} M_j^{-1} A_j$. Since M_j fixes $A_j W_j$, we see that $M_j A_j^{-1} M_j^{-1} A_j$ restricted to W_j is identical to M_j restricted to W_j.

In particular, $M_j A_j^{-1} M_j^{-1} A_j$ is a $\mathrm{P}(r)$-matrix, and it has degree at most $2\deg(A_j)$. Now we use Lemma 4.4 again, raising $M_j A_j^{-1} M_j^{-1} A_j$ to a large power, and we would obtain a matrix A_{j+1} of degree at most $\frac{2\deg(A_j)}{4}$, with eigenvalues either 1 or outside of $\mathbb{F}_q$.

Since $M_j \in T'$, by choosing T' as in Corollary 2.4, M_j have length bounded by $q^{2cn(\log n + \log q)^3}$. And since the order of $M_j A_j^{-1} M_j^{-1} A_j$ is bounded by q^n by Lemma 5.4, the length of A_{j+1} is at most

$$q^n (2q^{2cn(\log n + \log q)^3} + 2q^{2cn(\log n + \log q)^3 + j(n+2)}) \leq q^{2cn(\log n + \log q)^3 + (j+1)(n+2)}.$$

We repeat the above induction $\frac{\log n}{\log 2} - 1$ times, or stop early if we hit degree $2c(\log n + \log q)^3$. The last A_j we obtained is the desired matrix of small degree and small length. □

6. *t-Transversal Sets for Orthogonal Groups, Symplectic Groups, and Unitary Groups*

Let's fix some notation for the discussion of the following three sections. Let V be a non-degenerate formed space of dimension n over the field $\mathbb{F}_q$, with a non-degenerate quadratic form Q (the orthogonal case), non-degenerate alternating bilinear form B (the symplectic case), or non-degenerate Hermitian form B with field automorphism σ (the unitary case). In the orthogonal case, we shall let B be the symmetric bilinear form obtained by polarizing Q, i.e., $B(v, w) = Q(v + w) - Q(v) - Q(w)$. Let G be the group of isometries for V.

Definition 6.1.

(i) A vector $v \in V$ is ***singular*** if $B(v, v) = 0$ and (if applicable) $Q(v) = 0$.
(ii) A pair of singular vectors $v, w \in V$ is called a ***hyperbolic pair*** if $B(v, w) = 1$.
(iii) The subspace generated by a hyperbolic pair is a ***hyperbolic plane***.
(iv) A subspace W of V is ***anisotropic*** if it contains no singular vector other than 0.
(v) A subspace is ***totally singular*** if the form B and (if applicable) the quadratic form Q restricted to it is the zero form.
(vi) Given any subspace W of V, we define its ***orthogonal complement*** to be $W^\perp := \{v \in V : B(v, w) = 0 \text{ for all } w \in W\}$. Two subspaces U, W of V are ***orthogonal*** if they are in each other's orthogonal complimant. We denote this as $U \perp W$.
(vii) The ***radical*** of V is $V^\perp$.
(viii) A subspace W is ***radical-free*** if $W \cap V^\perp = \{0\}$.

Theorem 6.2 (Witt's Decomposition Theorem). *The non-degenerate formed space V has an orthogonal decomposition $V = V_{ani} \oplus (\bigoplus_{i=1}^{m} H_i)$, where V_{ani} is anisotropic of dimension at most 2, and H_i are hyperbolic planes. In particular, V has a totally singular subspace of dimension at least $\frac{\dim(V)-2}{2}$, and any anisotropic space in V has dimension at most 2.*

Proof. See [Gro02]. □

Lemma 6.3. *Recall that V is a non-degenerate formed space.*

(i) *$V^\perp = \{0\}$ unless the non-degenerate form for V is a quadratic form, and* $\operatorname{char} \mathbb{F}_q = 2$.
(ii) *$V^\perp$ has dimension at most 1.*
(iii) *For any subspace W,* $\dim W + \dim W^\perp$ *is equal to* $\dim V$ *if W is radical-free, and* $\dim V + 1$ *if W is not.*
(iv) *For any subspace W, $(W^\perp)^\perp = W + V^\perp$.*
(v) *A totally singular subspace is always radical-free.*

Proof. See [Gro02]. □

Definition 6.4. A subset S of G is called a ***singularly t-transversal set*** if, for any isometric embedding X of a t-dimensional totally singular subspace W into V, we can find $A \in S$ that extends X on W.

Lemma 6.5 (Witt's Extension Lemma). *G is a singularly t-transversal set for any t.*

Proof. This is a special case of Witt's extension lemma, which states that any bijective isometry of radical-free subspaces of V could be extended to an isometry of the whole formed space. See [Gro02] for a proof. □

Now, since our focus is on the finite simple groups, we don't really use the full isometry group G. Rather, we are interested in its commutator subgroup G'.

LEMMA 6.6. *For any $t \leq \frac{n-2}{5}$, the commutator subgroup G' of G is singularly t-transversal.*

Proof. Let W be a totally singular space of dimension t. Let $X : W \to V$ be any isometric embedding from W to V.

Step 1: I claim that there is a totally singular subspace W', which is orthogonal to W and $X(W)$, has trivial intersection with W and $X(W)$, and has the same dimension as W.

To see this, we have $\dim W^{\perp} = \dim X(W)^{\perp} \geq n - t$. Therefore, $\dim(W^{\perp} \cap X(W)^{\perp}) \geq n - 2t$. So in the subspace $W^{\perp} \cap X(W)^{\perp}$, we can find a subspace W'' of dimension $n - 3t$ with trivial intersections with W and $X(W)$. Now, since W'' is a formed space (possibly degenerate), it has a totally singular subspace of dimension at least $\frac{\dim W''-2}{2} = \frac{n-3t-2}{2} \geq t$. So, from this totally singular space, we could simply pick any totally singular subspace of dimension t to be the desired W'.

Step 2:

Let $Y : W \to W'$ be any bijective linear map. Since both spaces are totally singular, Y is an isometry. So we could find an extension $A \in G$.

Let $Z : W \oplus W' \to X(W) \oplus W'$ be the linear map that restricts to X on W, and restricts to the identity map on W'. Then by our choice of W', this is a well-defined isometry of totally singular subspaces, and it would have an extension $B \in G$.

Consider $BA^{-1}B^{-1}A \in G'$. This would restrict to X on W. So we are done. □

LEMMA 6.7. *Let S be any subset of G. If the subgroup generated by S is singularly t-transversal, then $\bigcup_{d=1}^{d=q^{nt}} S^d$ is singularly t-transversal.*

Proof. Let H be the subgroup generated by S. Let W be any t-dimensional totally singular subspace, and let $L(W)$ be the set of isometric embeddings of W into V. Then an element $g \in H$ acts on $L(W)$ by $g(X) = (g \circ X)|_W$ for any $X \in L(W)$. Let Γ be the corresponding Schreier graph of this action of H on $L(W)$ with generating set S.

Any isometric embedding from W to V is a linear map. Therefore, there are at most q^{nt} vertices for Γ, where $t = \dim W$. And since H is singularly t-transversal, the graph Γ must be connected. So Γ must have a diameter at most q^{nt}. □

COROLLARY 6.8. *Given any symmetric generating set S for G or G', the set $\bigcup_{d=1}^{d=q^{nt}} S^d$ is singularly t-transversal for $t \leq \frac{n-2}{5}$.*

7. *Degree Reducing for Orthogonal, Symplectic, Unitary Groups*

LEMMA 7.1. *For any non-zero singular $v \in V$, there is a vector $w \in V$ such that v, w form a hyperbolic pair.*

Proof. Recall that V is a non-degenerate formed space with an alternating bilinear, symmetric bilinear or Hermitian form B. In the case of characteristic 2, a symmetric bilinear form B might be degenerate, even though the formed space itself is not. Let σ be the field automorphism of the base field F for the Hermitian form B, or identity if B is bilinear.

For any element $k \in F$, we define $\mathrm{Tr}(x) = x + \sigma(x)$.

Now given a singular $v \in V$, since V is a non-degenerate formed space, we can find a vector $w' \in V$ such that $B(v, w') \neq 0$. By scaling w', we can assume that $B(v, w') = 1$.

Suppose we can find an element $k \in F$ such that $\mathrm{Tr}(k) = B(w', w')$, then $w = w' - kv$ is the desired vector forming a hyperbolic pair with v: $B(v, w) = B(v, w') = 1$, and

$$\begin{aligned} &B(w' - kv, w' - kv) \\ =&B(w', w') - \mathrm{Tr}(k) \\ =&0. \end{aligned}$$

Now, it remains to show that such k always exists.

Let E be the subfield of F fixed by σ. Then obviously $B(w', w') \in E$. So it is enough to show that either $E = \mathrm{Tr}(F)$, or $B(w', w') = 0$ for all w'.

Now, $\mathrm{Tr}(F)$ is closed under addition, and it is also closed under multiplication by elements of E. So $\mathrm{Tr}(F)$ is a E-vector space contained in E. So either $E = \mathrm{Tr}(F)$, or $\mathrm{Tr}(F) = 0$.

In the case that $\mathrm{Tr}(F) = 0$, then $\sigma(x) = -x$ for all $x \in F$. But since σ is a field automorphism, we must conclude that the field F has characteristic 2, and σ is the identity. Then the form B is alternating, and $B(w', w') = 0$ for any $w' \in V$. □

LEMMA 7.2. *If a subspace H of V is an orthogonal sum of hyperbolic planes, then $H \cap H^{\perp} = \{0\}$.*

Proof. The subspace H is an orthogonal sum of hyperbolic planes. Then let us assume that these planes are the linear span of hyperbolic pairs $(v_1, w_1), (v_2, w_2), (v_3, w_3), ..., (v_t, w_t)$.

Suppose $v \in H \cap H^{\perp}$. Then for some scalars $a_i, b_i \in F$, we have

$$v = \sum_{i=1}^{t} a_i v_i + \sum_{i=1}^{t} b_i w_i.$$

Now, since $B(v, v_i) = 0$, we can deduce that $b_i = 0$. Similarly, since $B(v, w_i) = 0$, we can deduce that $a_i = 0$. So $v = 0$. □

LEMMA 7.3. *Fix any nonzero elements $a, b, c \in \mathbb{F}_q$. Then the equation $ax^2 + by^2 + cz^2 = 0$ has a non-trivial solution in $\mathbb{F}_q$.*

Proof. If $\mathrm{char}(\mathbb{F}_q) = 2$, then $(\mathbb{F}_q)^*$ is a multiplicative group of odd order. So every nonzero element of $\mathbb{F}_q$ is a square.

Find x, y, z such that $x^2 = a^{-1}$, $y^2 = b^{-1}$ and $z = 0$. This is a non-trivial solution of the equation.

Suppose q is odd. Let S be the set of squares in $\mathbb{F}_q$. Then $|S| = \frac{q+1}{2}$. Then $|aS| + |-c - bS| > |\mathbb{F}_q|$. As a result, we have $aS \cap (-c - bS) \neq \varnothing$. So $-c \in aS + bS$.

Pick $x, y \in \mathbb{F}_q$ such that $ax^2 + by^2 = -c$. Then the triple $(x, y, 1)$ is a non-trivial solution to the equation. □

LEMMA 7.4. *Fix any $A \in G$. Then given any totally singular subspace $W \in V$ of dimension d, we can find a subspace W' of W such that W' is perperdicular to AW', and W' has dimension at least $\frac{d}{4} - \frac{3}{2}$.*

Proof. We proceed by induction on the dimension of W.

Initial Step: For the base case of the induction, suppose the dimension of W is 7 or 8 or 9 or 10. Then all we need is to find a nonzero vector $v \in W$ such that $v \perp Av$. Suppose for contradiction that there is no such vector.

Pick any non-zero $v_1 \in W$. Let W_1 be the intersection of W and $\text{span}\{v_1, Av_1, A^{-1}v_1\}^{\perp}$. Since v_1 is not perpendicular to Av_1, it is not in $\text{span}\{v_1, Av_1, A^{-1}v_1\}^{\perp}$. So $v_1 \notin W_1$, and W_1 has dimension at least $\dim W - 3 \geq 4$.

Pick any non-zero $v_2 \in W_1$. Let W_2 be the intersection of W_1 and $\text{span}\{v_2, Av_2, A^{-1}v_2\}^{\perp}$. Then W_2 has dimension at least $\dim W_1 - 3 \geq 1$ and similarly $v_2 \notin W_2$. Pick any non-zero $v_3 \in W_2$.

Now, we know $B(v_1, Av_1), B(v_2, Av_2), B(v_3, Av_3)$ are all in $\mathbb{F}_q^*$. We shall divide our discussion into two cases:

Orthogonal or Symplectic Case: Let $a = B(v_1, Av_1), b = B(v_2, Av_2), c = B(v_3, Av_3)$. Then by Lemma 7.3, we can find a nontrivial triple $x, y, z \in \mathbb{F}_q$ such that $ax^2 + by^2 + cz^2 = 0$. Let $v = xv_1 + yv_2 + zv_3$, then we have

$$\begin{aligned} B(v, Av) =& x^2 B(v_1, Av_1) + y^2 B(v_2, Av_2) + z^2 B(v_3, Av_3) \\ =& ax^2 + by^2 + cz^2 \\ =& 0. \end{aligned}$$

Unitary Case: If B is a Hermitian form for a field automorphism σ of order 2, then let F be the fixed subfield of σ. Let $N : \mathbb{F}_q \to F$ be the field norm, which is surjective.

Now, $\mathbb{F}_q$ is an F-vector space of dimension 2. So $B(v_1, Av_1), B(v_2, Av_2), B(v_3, Av_3)$ cannot be F-linearly independent in $\mathbb{F}_q$. So one can find non-trivial triple $a_1, a_2, a_3 \in F$ such that $a_1 B(v_1, Av_1) + a_2 B(v_2, Av_2) + a_3 B(v_3, Av_3) = 0$.

Since the norm map is surjective, find $x_1, x_2, x_3 \in \mathbb{F}_q$ such that $N(x_i) = a_i$. Let $v = x_1 v_1 + x_2 v_2 + x_3 v_3$. Then we have

$$\begin{aligned} B(v, Av) =& N(x_1) B(v_1, Av_1) + N(x_2) B(v_2, Av_2) + N(x_3) B(v_3, Av_3) \\ =& a_1 B(v_1, Av_1) + a_2 B(v_2, Av_2) + a_3 B(v_3, Av_3) \\ =& 0. \end{aligned}$$

So in either case, we could find the desired non-trivial vector $v \in W$ such that $v \perp Av$.

Inductive Step: Now let us proceed for general W of larger dimension. Since the dimension of W is at least 7, by the argument in the base case of the induction, we can find $v_1 \in W$ such that $B(v_1, Av_1) = 0$. Let W_1 be the intersection of W and $\text{span}\{v_1, Av_1, A^{-1}v_1\}^{\perp}$. Then W_1 has dimension at least $d - 3$. Pick any subspace W_2 of W_1 linearly independent from v_1. Then W_2 has dimension at least $d - 4$ and at most $d - 1$. Then by induction hypothesis, we can find W_2' a subspace of W_2, such that W_2' is perpendicular to AW_2', and W_2' has dimension at least $\frac{d-4}{4} - \frac{3}{2}$.

Let W' be the span of W_2' and v_1. Then W' will be perpendicular to AW', and has dimension at least $\frac{d-4}{4} - \frac{3}{2} + 1 = \frac{d}{4} - \frac{3}{2}$. So we are done. □

LEMMA 7.5. *Fix any $A \in G$ where all eigenvalues of A are outside of $\mathbb{F}_q$. Then there is a t-dimensional totally singular subspace W of V such that $W \cap AW = \{0\}$, for any $t \leq \frac{n}{6}$.*

Proof.

Fix n, which we assume to be at least 3, so that V has at least one singular vector. We shall proceed by induction on the dimension of W.

Initial Step: Suppose $t = 1$. Simply pick any singular vector v, and let W be the span of v. Since A has no eigenvalue in $\mathbb{F}_q$, v and Av must be linearly independent. So $W \cap AW = \{0\}$.

Inductive Step: Suppose we have found a totally singular subspace W of dimension $t-1$ such that $W \cap AW = \{0\}$. I claim that, when $t \leq \frac{n}{6}$, we can find another singular vector v, such that the desired subspace is the span of v and W.

First of all, we want v to be a singular vector perpendicular to W. We know $W^\perp$ has dimension $n-t+1$, and by Witt's decomposition theorem, V has a totally singular space of dimension at least $\frac{n-2}{2}$. This totally singular space will intersect $W^\perp$ in a subspace of dimension at least $\frac{n-2}{2} - t + 1 = \frac{n}{2} - t$. So there are at least $q^{\frac{n}{2}-t}$ singular vectors perpendicular to W.

Among these vectors, to prove the existence of a good v, we should count the number of vectors to avoid. We need v to avoid $W + AW$. Afterwards, it is enough to have Av avoiding any linear combination of v and $W + AW$. To satisfy the second requirement, we need v to avoid $\bigcup_{x\in\mathbb{F}_q}(A-x)^{-1}(W+AW)$. Here we shall interpret $(A-x)^{-1}$ as the pullback map of subsets.

Now, since A has no eigenvalue in $\mathbb{F}_q$, therefore $A - x$ are all invertible. So, we have

$$\begin{aligned}&|(W+AW) \cup (\bigcup_{x\in\mathbb{F}_q}(A-x)^{-1}(W+AW))|\\ \leq& q^{2t-2} + q \times q^{2t-2}\\ <& q^{2t}.\end{aligned}$$

So as long as $2t \leq \frac{n}{2} - t$, i.e., $t \leq \frac{n}{6}$, then it is possible to choose a vector v as desired. □

LEMMA 7.6. *Fix any matrix $A \in G$ of degree k, such that the eigenvalues of A are either 1 or outside of $\mathbb{F}_q$. Then we can find a subspace W of V with the following properties:*

(i) $\dim W \geq \frac{k}{32} - \frac{7}{4}$
(ii) *W is totally singular;*
(iii) $W \cap AW = \{0\}$*;*
(iv) $W \perp AW$.

Proof. Let V_A be the subspace of fixed points of A in V. Let $V_r = V_A \cap (V_A)^\perp$. Choose any positive number a to be determined later. Then either V_r has dimension $< a$, or it has dimension $\geq a$.

Case of Large V_r:

Suppose V_r has dimension $\geq a$. Pick any non-zero singular $v_1 \in V_r$, then we can find $w_1 \in V$ such that v_1, w_1 form a hyperbolic pair. Let V_{r1} be the intersection of V_r with $\mathrm{span}\{v_1, w_1\}^\perp$. Pick any non-zero singular $v_2 \in V_{r1}$, then we can find w_2 in $\mathrm{span}\{v_1, w_1\}^\perp$, such that v_2, w_2 form a hyperbolic pair. Then let V_{r2} be the intersection of V_{r1} with $\mathrm{span}\{v_1, w_1, v_2, w_2\}^\perp$, and repeat.

As long as $\dim V_{ri} > 2$, then V_{ri} cannot be anisotropic. So we can keep going at least $\lfloor \frac{a-2}{2} \rfloor$ times. Thus we obtained $w_1, ..., w_{\lfloor \frac{a-2}{2} \rfloor}$. They span a totally singular space W_r of dimension at least $\frac{a-3}{2}$. Then by Lemma 7.4, we can find a subspace W of W_r, such that $W \perp AW$ and W has dimension at least $\frac{a-3}{8} - \frac{3}{2}$.

I claim that, ignoring the dimension requirement, this W satisfies all the desired properties. By construction of W, we have W totally singular and $W \perp AW$. We only need to show that $W \cap AW = \{0\}$.

For any vector $w = \sum_{i=1}^{\lfloor \frac{a-2}{2} \rfloor} a_i w_i \in W$, suppose it is perpendicular to V_r. Then for each i, since $B(v_i, w) = 0$, we see that $a_i = 0$. So $w = 0$. To sum up, W has trivial intersection with $(V_r)^\perp$.

Suppose $w \in W \cap AW$. Then $w - A^{-1}w \in W$, and for any $v \in V_r$ we have

$$\begin{aligned} B(v, w - A^{-1}w) =& B(v, w) - B(v, A^{-1}w) \\ =& B(v, w) - B(Av, w) \\ =& B(v, w) - B(v, w) \\ =& 0. \end{aligned}$$

So $w - A^{-1}w \in W \cap V_r^\perp = \{0\}$. So $w = Aw$, and $w \in W \cap V_A \subseteq W \cap (V_r)^\perp = \{0\}$.

To sum up, this W is the space we desired, with dimension at least $\frac{a-3}{8} - \frac{3}{2}$.

Case of Small V_r:

Suppose V_r has dimension $< a$.

Step 1: We want to first find a subspace W_A of $(V_A)^\perp$ where $W_A \perp W_A$, and $W_A \cap AW_A = V_r$, and the codimension of V_r in W_A is at least $\frac{k-a-5}{6}$.

Now, the bilinear or Hermitian form B restricted to $(V_A)^\perp$ is still bilinear or Hermitian, with exactly V_r as the radical. So the space $V' = (V_A)^\perp / V_r$ has an induced bilinear or Hermitian form B', and now B' is non-degenerate.

So V' is a non-degenerate formed space with dimension at least $k - a$. Furthermore, since V_r and $(V_A)^\perp$ are both A-invariant, A induces a linear map A' on V'. Clearly A' has no non-trivial fixed point in V', so all eigenvalues of A' are ouside of $\mathbb{F}_q$. So by Lemma 7.5, V' has a totally singular subspace W' of dimension at least $\lfloor \frac{k-a}{6} \rfloor \geq \frac{k-a-5}{6}$, such that $W' \cap A'W' = \{0\}$.

Let W_A be the pullback of W' through the projection map $(V_A)^\perp \to V'$. Since W' is totally singular under B', the form B vanishes on W_A. (Note that in the orthogonal case, the quadratic form Q might not vanish on W_A, so W_A might not be totally singular.)

Step 2: Now let us find a totally singular subspace W_r of W_A intersecting trivially with V_r and has dimension at least $\frac{k-a-5}{6}$.

If $\operatorname{char} \mathbb{F}_q \neq 2$, or if we are not in the orthogonal case, then W_A is totally singular. Pick any subspace W_r of W_A which has trivial intersection with V_r and has dimension at least $\frac{k-a-5}{6}$, and we are done.

Suppose $\operatorname{char} \mathbb{F}_q = 2$, and we are in the orthogonal case, and Q vanishes on V_r. Then the space $V' = (V_A)^\perp / V_r$ would have an induced non-degenerate quadratic form Q' that corresponds to the non-degenerate bilinear form B'. So by Lemma 7.5, when we picked W' to be totally singular, we can pick it to be totally singular with respect to the quadratic form Q'. This way the subspace W_A would be totally singular. Again pick any subspace W_r of W_A which has trivial intersection with V_r and has dimension at least $\frac{k-a-5}{6}$, and we are done.

Finally, suppose now that $\operatorname{char} \mathbb{F}_q = 2$, and we are in the orthogonal case, and we have a vector $v_0 \in V_r$ such that $Q(v_0) \neq 0$. Since $\operatorname{char} \mathbb{F}_q = 2$, the squaring map is bijective on $\mathbb{F}_q$, we can assume that $Q(v_0) = 1$ by scaling v_0.

Define a map $X : W_A \to W_A$ such that $X(v) = v + \sqrt{Q(v)} v_0$. Here the square root is well defined because $\operatorname{char} \mathbb{F}_q = 2$. Then we have $Q(X(v)) = 0$ for all $v \in W_A$.

Furthermore, X is linear. To see this, first we notice that for any v, w in W_A, since B vanishes on W_A,

$$Q(v + w) = Q(v) + Q(w) + B(v, w) = Q(v) + Q(w).$$

So we have

$$\begin{aligned} X(v+w) =& v+w+\sqrt{Q(v+w)}v_0 \\ =& v+w+\sqrt{Q(v)+Q(w)}v_0 \\ =& v+w+(\sqrt{Q(v)}+\sqrt{Q(w)})v_0 \\ =& X(v)+X(w). \end{aligned}$$

For any scalar $a \in \mathbb{F}_q$, we also easily have $X(av) = aX(v)$.

Now, since X is linear, $X(W_A)$ is a subspace of W_A. So $X(W_A)$ is a totally singular subspace.

Now pick any subspace W_r of W_A which has trivial intersection with V_r and has dimension at least $\frac{k-a-5}{6}$. Then $X(W_r)$ is a totally singular subspace of W_A. It remains to show that this $X(W_r)$ intersects V_r trivially and has the correct dimension.

For any vector v, if $X(v) \in V_r$, then $v = \sqrt{Q(v)}v_0 + X(v) \in V_r$. So $X(W_r)$ only has trivial intersection with V_r. And since the kernel of X is entirely in V_r, $X(W_r)$ has the same dimension as W_r.

So replace W_r by $X(W_r)$, and we are done.

Step 3: Now we construct the desired subspace W.

By Lemma 7.4, we find a subspace W of W_r such that $W \perp AW$, and W has dimension at least $\frac{k-a-5}{24} - \frac{3}{2}$.

I claim that, ignoring the dimension requirement, this W satisfies all the desired properties.

First of all, we know W is totally singular and $W \perp AW$. By construction, W is in $(V_A)^\perp$ but has trivial intersection with V_r. Then since $W_A \cap AW_A = V_r$, we know that $W \cap AW = \{0\}$.

To sum up, this W is the space we desired, with dimension at least $\frac{k-a-5}{24} - \frac{3}{2}$.

Find the Optimal a:

Picking the optimal $a = \frac{k}{4} + 1$ for both cases above, we eventually find the desired subspace W of dimension at least $\frac{k}{32} - \frac{7}{4}$. □

PROPOSITION 7.7. *For any symmetric generating set S of G or G', there is a non-trivial element of degree at most $C(\log n + \log q)^3$ for some absolute constant C, of length less than $q^{C' n (\log n + \log q)^3}$ for some absolute constant C'.*

Proof. We pick r and c according to Lemma 4.4. Let us assume that $2c(\log n + \log q)^3 < \frac{n-2}{5}$, because otherwise the statement is trivial.

Let $p_1, ..., p_r$ be the first r primes coprime to $p(q-1)$. Let $f_i(x)$ be the irreducible polynomial over $\mathbb{F}_q$ for all the primitive p_i-th roots of unity, and let C_i be the companion matrix of $\mathbb{F}_q$.

Initial Step:

Let us find our first P(r)-matrix. Let T be a singularly $c(\log n + \log q)^3$-transversal set. Let W be any totally singular subspace of dimension $c(\log n + \log q)^3$, which exists by Witt's decomposition theorem. Note that any bijective linear map from W to W is an isometry, and is therefore subject to Witt's extension lemma.

By definition of a singularly transversal set, we can find $A_0 \in T$ that maps the totally singular subspace W onto itself, and that its restriction to this subspace is the matrix $(\bigoplus_{i=1}^{r} C_i) \oplus I$ for some arbitrary choices of basis on W, where I is some identity matrix of suitable size.

In particular, A_0 is a P(r)-matrix. Since $A_0 \in T$, by choosing T as in Corollary 6.8, A_0 has length bounded by $q^{cn(\log n + \log q)^3}$.

By using Lemma 4.4, we can raise A_0 to a large power, and obtain a non-identity matrix A_1 of degree $\leq \frac{\deg(A_0)}{4} \leq \frac{n}{4}$, with eigenvalues either 1 or outside of $\mathbb{F}_q$. Since the order of A_0 is bounded by q^n by Lemma 5.4, the length of A_1 is bounded by $q^{cn(\log n + \log q)^3 + n}$.

Inductive Step:

Suppose we have obtained a non-identity matrix A_j with eigenvalues either 1 or outside of $\mathbb{F}_q$, degree at most $\frac{n}{2^{j+1}}$, and length at most $q^{2cn(\log n+\log q)^3+j(n+2)}$. If $\deg A_j \leq 56+32c(\log n+\log q)^3$, then we stop. If not, then let us construct a non-identity matrix A_{j+1} of even smaller degree.

First we shall transform A_j into a $\mathrm{P}(r)$-matrix. Find a totally singular subspace W_j of dimension $c(\log n+\log q)^3$ as in Lemma 7.6. In particular, $W_j \oplus A_jW_j$ is a well-defined totally singular space. Let T' be a singularly $2c(\log n+\log q)^3$-transversal set, then we can find $M_j \in T'$ that fixes A_jW_j, and restricts to a map from W_j to W_j as $\bigoplus_{i=1}^{r} C_i \oplus I$ for any arbitrary basis of W_j and some identity matrix I of suitable size.

Consider the commutator $M_jA_j^{-1}M_j^{-1}A_j$. Since M_j fixes A_jW_j, we see that $M_jA_j^{-1}M_j^{-1}A_j$ restricted to W_j is identical to M_j restricted to W_j.

In particular, $M_jA_j^{-1}M_j^{-1}A_j$ is a $\mathrm{P}(r)$-matrix, and it has degree at most $2\deg(A_j)$. Now we use Lemma 4.4 again, raising $M_jA_j^{-1}M_j^{-1}A_j$ to a large power, and we would obtain a matrix A_{j+1} of degree at most $\frac{2\deg(A_j)}{4}$, with eigenvalue wither 1 or outside of $\mathbb{F}_q$.

Since $M_j \in T'$, by choosing T' as in Corollary 6.8, M_j have length bounded by $q^{2cn(\log n+\log q)^3}$. And since the order of $M_jA_j^{-1}M_j^{-1}A_j$ is bounded by q^n by Lemma 5.4, the length of A_{j+1} is at most

$$q^n(2q^{2cn(\log n+\log q)^3}+2q^{2cn(\log n+\log q)^3+j(n+2)}) \leq q^{2cn(\log n+\log q)^3+(j+1)(n+2)}.$$

We repeat the above induction $\frac{\log n}{\log 2}-1$ times, or stop early if we hit degree $2c(\log n+\log q)^3$. The last A_j we obtained is the desired matrix of small degree and small length. □

8. *The Conjugacy Expansion Lemmas*

In this section, we shall show that any small degree element will quickly generate the whole group with any symmetric generating set.

LEMMA 8.1. *Let S be any symmetric generating set for a subgroup H of $\mathrm{GL}_n(\mathbb{F}_q)$. Let A be any matrix in H of degree k, and let B be any matrix conjugate to A in H. Then $B = MAM^{-1}$ for some $M \in H$ of length at most q^{2nk}.*

Proof. Since A has degree k, we know $A = I + A'$ for some matrix A' of rank k. So we can decompose A' as a product XY where X is an n by k matrix of full rank and Y is a k by n matrix of full rank. So $A = I + XY$.

Any conjugates of A can similarly be expressed as $I + X'Y'$ where X' is some n by k matrix of full rank, and Y' is some k by n matrix of full rank. There are at most q^{2nk} possibilities for the pair (X', Y'). So there are at most q^{2nk} conjugates of A.

Now H acts on the conjugacy class of A in H by left conjugation, and the corresponding Schreier graph must be connected. So the Schreier graph has diameter bounded by the number of vertices, i.e., q^{2nk}. □

THEOREM 8.2 ([LS01]). *Let G be $\mathrm{SL}_n(\mathbb{F}_q)$, $\Omega_n(\mathbb{F}_q)$, $\mathrm{Sp}_n(\mathbb{F}_q)$, or $\mathrm{SU}_n(\mathbb{F}_q)$. Let $A \in G$ be an element of degree k outside the center of G. Then every element of G is a product of at most $O(\frac{n}{k})$ conjugates of A.*

PROPOSITION 8.3. *Let G be $\mathrm{SL}_n(\mathbb{F}_q)$, $\Omega_n(\mathbb{F}_q)$, $\mathrm{Sp}_n(\mathbb{F}_q)$, or $\mathrm{SU}_n(\mathbb{F}_q)$. Let S be any symmetric generating set for G. Suppose we have a non-trivial element $A \in G$ of length $d > 0$ and degree $k < n$. Then the diameter of G with respect to S will be $O((2q^{2nk}+d)\frac{n}{k})$.*

Proof. For any B conjugate to A, by the Lemma 8.1 above, $B = MAM^{-1}$ for some $M \in G$ of length at most q^{2nk}. So B has length at most $2q^{2nk} + d$. So every conjugate of A in G has length bounded by $2q^{2nk} + d$.

Furthermore, since A has degree $< n$ but non-trivial, it is not a scalar matrix. Then by [Gro02], A is not in the center of G.

Now by the resulf of Liebeck and Shalev [LS01], we know that every element of G can be written as a product of $O(\frac{n}{k})$ conjugates of A. So the whole group G has a diameter bound of $O((2q^{2nk} + d)\frac{n}{k})$. □

COROLLARY 8.4. *The diameter of a finite simple group of Lie type of rank n over $\mathbb{F}_q$ are at most $O(q^{O(n(\log n+\log q)^3)})$, independent of the choice of generating sets. The implied constants are absolute.*

Proof. Combine Proposition 8.3 with Proposition 5.5 or Proposition 7.7. □

9. *Implications on Spectral Gap and Mixing Time*

Given a group G and its generating set S. Let $\Gamma(G, S)$ be its Cayley graph, and let A be the normalized adjacency matrix of the graph. Then A has real eigenvalues $\lambda_1, ..., \lambda_{|G|}$, ordered from the largest one to the smallest one. Then the ***spectral gap*** of $\Gamma(G, S)$ is $\lambda_1 - \lambda_2$.

Let μ be the random distribution $\frac{1}{2}1_{\{e\}} + \frac{1}{2|S|}1_S$. Then a lazy random walk of length k is the random outcome of the distribution $\mu^{(k)} = \mu * \mu * \mu * ... * \mu$. Using the definition of Helfgott, Seress and Zuk [HSZ15], the ***strong mixing time*** of $\Gamma(G, S)$ is the least number k such that $\mu^{(k)}$ is at most $\frac{1}{2|G|}$ away from the uniform distribution on $\Gamma(G, S)$, in the ℓ^∞ norm.

One can bound the spectral gap using a diameter bound.

PROPOSITION 9.1 ([DSC93], Corollary 3.1). *Given a finite group G and a symmetric generating set S, let Γ be the Cayley graph. Then the spectral gap of the Cayley graph is bounded from below by $\frac{1}{(\operatorname{diam}\Gamma)^2|S|}$*

In turn, one can bound the strong mixing time by the spectral gap.

PROPOSITION 9.2 ([Lovsz], Theorem 5.1). *Given a finite group G and a symmetric generating set S, let Γ be the Cayley graph, and let λ be the spectral gap. Then the strong mixing time of the Cayley graph is bounded by $O(\frac{\log|\Gamma|}{\lambda})$.*

Then our main result implies the following corollary:

COROLLARY 9.3. *Let G be a finite simple group of Lie type of rank n over $\mathbb{F}_q$. The spectral gap of $\Gamma(G, S)$ is bounded by $|S|^{-1}q^{-O(n(\log n+\log q)^3)}$, and the mixing time of $\Gamma(G, S)$ is bounded by $|S|q^{O(n(\log n+\log q)^3)}$.*

References

[BS88] L. Babai and A. Seress. "On the diameter of Cayley graphs of the symmetric group". ***J. Combin. Theory***. A 49(1) (1988), pp. 175–179.

[BS92] L. Babai and A. Seress. "On the diameter of permutation groups". ***European J. Combin.*** 13(4) (1992), pp. 231–243.

[BGT11a] E. Breuillard, B. Green, and T. Tao. "Approximate subgroups of linear groups". ***Geom. Funct. Anal.*** 21(4) (2011), pp. 774–819.

[BT16] E. Breuillard and M. Tointon. "Nilprogressions and groups with moderate growth". ***Adv. Math.*** 289 (2016), pp. 1008–1055.

[BGT11b] E. Breuillard, B. Green, and T. Tao. "Suzuki groups as expanders". ***Groups Geom. Dyn.*** 5 (2011).

[Bre+15] E. Breuillard et al. "Expansion in finite simple groups of Lie type". ***J. Eur. Math. Soc. (JEMS)*** 17 (2015).

[DSC93] P. Diaconis and L. Saloff-Coste. "Comparison techniques for random walk on finite groups". ***Ann. Probab.*** 21(4) (1993), pp. 2131–2156.

[Fou85] E. Fouvry. "The Brun-Titchmarsh theorem: application to the Fermat theorem". ***Invent. Math.*** 79(2) (1985), pp. 383–407.

[Gro02] L. C. Grove. ***Classical groups and geometric algebra***. American Mathematical Soc., 2002.

[Hel08] H. A. Helfgott. "Growth and generation in $SL_2(\mathbb{Z}/p\mathbb{Z})$". ***Ann. of Math. (2)*** 167(2) (2008), pp. 601–623.

[HS14] H. A. Helfgott and A. Seress. "On the diameter of permutation groups". ***Ann. of Math. (2)*** 179(2) (2014), pp. 611–658.

[HSZ15] H. A. Helfgott, A. Seress, and A. Zuk. "Random generators of the symmetric group: diameter, mixing time and spectral gap". ***J. Algebra*** 421 (2015), pp. 349–368.

[Her75] I. N. Herstein. ***Topics in Algebra 2nd Edition***. Xerox College Publishing, 1975.

[Lan02] S. Lang. ***Algebra revised third edition***. GTM 211, 2002.

[LS01] M. W. Liebeck and A. Shalev. "Diameters of finite simple groups: sharp bounds and applications". ***Ann. of Math. (2)*** 154(2) (2001), pp. 383–406.

[Lovsz] L. Lovász. "Random walks on graphs: a survey". ***Combinatorics, Paul Erdős is eighty*** 2 (Keszthely, 1993), pp. 353–397.

[PS16] L. Pyber and E. Szabó. "Growth in finite simple groups of Lie type". ***J. Amer. Math. Soc.*** 29(1) (2016), pp. 95–146.